arXiv:1601.01886v2 [math.CO] 2 Oct 2016

Pathwidth and nonrepetitive list coloring

Adam Gągol
Theoretical Computer Science Department
Faculty of Mathematics and Computer Science
Jagiellonian University
Kraków, Poland
gagol@tcs.uj.edu.pl

Gwenaël Joret*
Computer Science Department
Université Libre de Bruxelles
Brussels, Belgium
gjoret@ulb.ac.be

Jakub Kozik† Piotr Micek†
Theoretical Computer Science Department
Faculty of Mathematics and Computer Science
Jagiellonian University
Kraków, Poland
{jkozik,micek}@tcs.uj.edu.pl

October 4, 2016

Abstract

A vertex coloring of a graph is *nonrepetitive* if there is no path in the graph whose first half receives the same sequence of colors as the second half. While every tree can be nonrepetitively colored with a bounded number of colors (4 colors is enough), Fiorenzi, Ochem, Ossona de Mendez, and Zhu recently showed that this does not extend to the list version of the problem, that is, for every $\ell \geqslant 1$ there is a tree that is not nonrepetitively ℓ-choosable. In this paper we prove the following positive result, which complements the result of Fiorenzi *et al.*: There exists a function f such that every tree of pathwidth k is nonrepetitively $f(k)$-choosable. We also show that such a property is specific to trees by constructing a family of pathwidth-2 graphs that are not nonrepetitively ℓ-choosable for any fixed ℓ.

1 Introduction

A *repetition* of length r ($r \geqslant 1$) in a sequence of symbols is a subsequence of consecutive terms of the form $x_1 \dots x_r x_1 \dots x_r$. A sequence is *nonrepetitive* (or *square-free*) if it does not contain a repetition of any length. In 1906 Thue proved that there exist arbitrarily long

*G. Joret was supported by a DECRA Fellowship from the Australian Research Council.
†J. Kozik and P. Micek were supported by the Polish National Science Center, grant 2011/01/D/ST1/04412.

nonrepetitive sequences over an alphabet of size 3 (see [2, 16]). The method discovered by Thue is constructive and uses substitutions over a given set of symbols.

A different approach to creating long nonrepetitive sequences was recently introduced by Grytczuk, Kozik, and Micek [10]: Generate a sequence by iteratively appending a random symbol at the end, and each time a repetition appears erase the repeated block. (For instance, if the sequence generated so far is *abcb* and we add *c*, then we erase the last two symbols, bringing us back to *abc*.) By a simple counting argument one can prove that with positive probability the length of the constructed sequence eventually exceeds any finite bound, provided the alphabet has size at least 4. This is one more than in Thue's result but the proof is more flexible and can be adapted to other settings. For instance, it led to a very short proof that for every $n \geqslant 1$ and every sequence of sets $L_1, \ldots, L_n$, each of size at least 4, there exists a nonrepetitive sequence $s_1 s_2 \ldots s_n$ where $s_i \in L_i$ for all i (see [10]), a theorem first proved by Grytczuk, Przybyło, and Zhu [11] via an intricate application of the Lefthanded Local Lemma. Whether the analogous statement for lists of size 3 is true remains an exciting open problem.

In this paper we make use of the above-mentioned approach to color trees nonrepetitively. Given an (undirected, simple) graph G, we denote by $V(G)$ and $E(G)$ its vertex set and edge set, respectively. A coloring $\phi : V(G) \to \mathbb{N}$ of the vertices of G is *nonrepetitive* if there is no repetition in the color sequence of any path in G; that is, ϕ is nonrepetitive if for every path P with an even number of vertices the sequence of colors on the first half of P is distinct from the sequence of colors on the second half of P. (We remark that all paths in this paper are simple, that is, contain no repeated vertex.) The minimum number of colors used in a nonrepetitive coloring of G is called the *Thue chromatic number* of G and is denoted by $\pi(G)$. Now, given a graph G, suppose that each vertex $v \in V(G)$ has a preassigned list of available colors $L_v \subset \mathbb{N}$. A coloring of G with these lists is a coloring ϕ of G such that $\phi(v) \in L_v$ for each vertex $v \in V(G)$. The *Thue choice number* of G, denoted by $\pi_l(G)$, is the minimum ℓ such that, for every list assignment $\{L_v\}_{v \in V(G)}$ with $|L_v| \geqslant \ell$ for each $v \in V(G)$, there is a nonrepetitive coloring of G with these lists.

Similarly as for many graph coloring parameters, the Thue chromatic (choice) number can be bounded from above by a function of the maximum degree: Alon, Grytczuk, Hałuszczak, and Riordan [1] proved that for every graph G with maximum degree Δ we have $\pi(G) \leqslant \pi_l(G) \leqslant c \cdot \Delta^2$ for some absolute constant c. A number of subsequent works [6, 8, 9, 12] focused on reducing the value of the constant c, the current best bound being $\pi_l(G) \leqslant (1 + o(1))\Delta^2$ (see [6]). Alon *et al.* [1] also showed that there are graphs with maximum degree Δ with $\pi(G) = \Omega\left(\frac{\Delta^2}{\log \Delta}\right)$. (Whether this can be improved by a $\log \Delta$ factor remains an open problem.)

It is not difficult to show that every tree has Thue chromatic number at most 4 (see [3]), which is best possible. This result was generalized to graphs of bounded treewidth by Kündgen and Pelsmajer [14]. They proved that $\pi(G) \leqslant 4^k$ for every graph G of treewidth k. It is not known whether this upper bound can be improved to a polynomial in k. However, if one considers graphs of pathwidth k instead, a polynomial bound is known: It was shown by Dujmović *et al.* [6] that $\pi(G) \leqslant 2k^2 + 6k + 1$ for every graph G of pathwidth k. (We note that quadratic might not be the right order of magnitude here.)

Probably the most intriguing open problem regarding the Thue chromatic number is whether it is bounded for all planar graphs, a question originally asked by Grytczuk [9]. A $O(\log n)$ upper bound is known [5], and from below Ochem constructed a planar graph requiring 11 colors (see [5]).

The main focus of this paper is the list version of the parameter, the Thue choice number. As mentioned at the beginning of the introduction, we have $\pi_l(P) \leqslant 4$ for every path P, and it is open whether this bound can be improved to 3. Fiorenzi, Ochem, Ossona de Mendez, and Zhu [7] gave the first example of a class of graphs where the Thue chromatic and Thue choice numbers behave very differently: While trees have Thue chromatic number at most 4, they showed that the Thue choice number of trees is unbounded. Clearly, trees with large Thue choice number must have large maximum degree, and in fact one can deduce from the proof in [7] that there are trees with maximum degree Δ and Thue choice number $\Omega(\frac{\log \Delta}{\log \log \Delta})$. Kozik and Micek [13] subsequently showed that a better-than-quadratic upper bound in terms of the maximum degree exists for trees: For every $\varepsilon > 0$ there exists $c > 0$ such that $\pi_l(T) \leqslant c\Delta^{1+\varepsilon}$ for every tree T of maximum degree Δ. (Bridging the significant gap between the upper and lower bounds remains an open problem.)

Note that graphs of bounded treewidth have unbounded Thue choice number since this is already the case for trees. On the other hand, Dujmović *et al.* [6] observed that $\pi_l(G)$ is bounded when G is a graph of pathwidth 1. This prompted the authors of [6] to ask whether $\pi_l(G)$ is bounded more generally when G has bounded pathwidth (which is the case for the Thue chromatic number). Also, since connected graphs G of pathwidth 1 are caterpillars, and thus trees in particular, they also asked the same question but with G moreover required to be a tree. A second motivation for the latter question was that the trees with arbitrarily large Thue choice number constructed by Fiorenzi *et al.* [7] also have unbounded pathwidth.

In this paper we answer both questions. First, we give a simple construction showing that the Thue choice number is unbounded for graphs of bounded pathwidth; in fact, this is true even for graphs of pathwidth 2 (which is best possible as noted above):

Theorem 1. *For every $\ell \geqslant 1$, there is a graph G of pathwidth 2 with $\pi_l(G) \geqslant \ell$.*

Next, we address the case of trees and prove that their Thue choice number is bounded from above by a function of their pathwidth:

Theorem 2. *There is a function $b : \mathbb{N} \to \mathbb{N}$ such that $\pi_l(T) \leqslant b(k)$ for every tree T of pathwidth k.*

The proof of Theorem 2 combines an induction on the pathwidth with the algorithmic method of Grytczuk *et al.* [10] to produce arbitrarily long nonrepetitive sequences described at the beginning of the introduction. This method, which finds its roots in the celebrated algorithmic proof of the Local Lemma by Moser and Tardos [15], was extended to produce nonrepetitive colorings of graphs (in [6]) and trees (in [13]). Part of our proof consists in adapting the ideas from [6, 13] to the situation under consideration.

We note that the bounding function $b(k)$ in Theorem 2 stemming from our proof is quite large, it is doubly exponential in k.

The paper is organized as follows: In Section 2 we introduce definitions and terminology. Then we prove Theorem 1 in Section 3, and Theorem 2 in Section 4.

2 Definitions

For an integer $n \geqslant 1$, we let $[n] := \{1, \ldots, n\}$. Also, given two integers a, b with $a \leqslant b$ we let $[a, b] := \{a, a+1, \ldots, b\}$, which we call an *interval*.

Graphs in this paper are finite, simple, and undirected. The vertex set and edge set of a graph G are denoted $V(G)$ and $E(G)$, respectively. Note that, since only simple graphs are considered, resulting loops and parallel edges are removed when contracting edges in a graph. A graph H is a *minor* of a graph G if H can be obtained from a subgraph of G by contracting edges.

A *tree decomposition* of a graph G is a pair $(T, \mathcal{C})$ where T is a tree and $\mathcal{C}$ is a collection $\{T_v : v \in V(G)\}$ of non-empty subtrees of T such that $V(T_u) \cap V(T_v) \neq \emptyset$ for every edge $uv \in E(G)$. The *width* of the tree decomposition $(T, \mathcal{C})$ is the maximum, over every $x \in V(T)$, of the number of subtrees in $\mathcal{C}$ containing x, minus 1. The *treewidth* of G is the minimum width of a tree decomposition of G. *Path decompositions* and *pathwidth* are defined analogously with the tree T required instead to be a path. Treewidth and pathwidth are minor-closed parameters, in the sense that every minor of a graph G has treewidth (pathwidth) at most that of G. We refer the reader to Diestel's textbook [4] for an introduction to the theory of treewidth and graph minors.

The *length* of a path is the number of its edges. The *height* of a rooted tree T is the maximum length of a path from the root to a leaf of T. (Thus T has height 0 if it consists of a unique vertex.) The *height* of a vertex v of T is the length of the path from the root to v in T.

3 Graphs of pathwidth 2

Let $G_{n,\ell}$ be the graph constructed from the path on $2n$ vertices where every second vertex is blown up to $\binom{\ell n}{\ell}$ vertices forming an independent set. Formally,

$$V(G_{n,\ell}) = \{v_{2i-1} \mid i \in [n]\} \cup \left\{ v_{2i}^j \mid i \in [n],\ j \in \left[\binom{\ell n}{\ell}\right] \right\},$$

and two vertices are adjacent in $G_{n,\ell}$ if and only if their lower indices differ by exactly 1. Also, let $V_i := \{v_i\}$ for each odd index $i \in [2n]$ and $V_i := \left\{v_i^j \mid j \in \left[\binom{\ell n}{\ell}\right]\right\}$ for each even index $i \in [2n]$.

It is not difficult to check that $G_{n,\ell}$ has pathwidth at most 2 (with equality for $n \geqslant 2$ and $\ell \geqslant 1$). Thus Theorem 1 follows from the following theorem.

Theorem 3. *Let ℓ and n be integers such that $\ell \geqslant 1$ and $n > e^{\ell+2}$. Then $\pi_l(G_{n,\ell}) > \ell$.*

Proof. Consider the following list assignment for the vertices of $G_{n,\ell}$. For each odd index $i = 2t + 1 \in [2n]$, vertex v_i is assigned the list

$$L_i := \{t\ell + 1, t\ell + 2, \ldots, t\ell + \ell\}.$$

Thus these n lists have size ℓ, are pairwise disjoint, and their union is $[\ell n]$. Next, enumerate the ℓ-subsets of $[\ell n]$ in an arbitrary way. Then, for each even index $i \in [2n]$ and index $j \in \left[\binom{\ell n}{\ell}\right]$, vertex v_i^j is assigned the list which is the j-th set in that enumeration.

We claim that, because n was chosen to be strictly larger than $e^{\ell+2}$, there cannot be a nonrepetitive coloring of $G_{n,\ell}$ with these lists. Arguing by contradiction, let us suppose that ϕ is such a coloring.

With a slight abuse of notation, for $i \in [2n]$ we use the shorthand $\phi(V_i)$ for the set $\cup_{u \in V_i} \{\phi(u)\}$. Consider an interval $I \subseteq [2n]$ of the form $I = [a, a + 4k + 1]$ with $k \geqslant 0$. Suppose that the following two conditions are satisfied:

$$\begin{aligned} &\phi(v_i) \in \phi(V_{i+2k+1}) \quad \text{for each } i \in [a, a + 2k], i \text{ odd} \\ &\phi(v_{i+2k+1}) \in \phi(V_i) \quad \text{for each } i \in [a, a + 2k], i \text{ even.} \end{aligned}$$

Then it is easy to check that there exists a path $w_a, \ldots, w_{a+4k+1}$ in $G_{n,\ell}$ with $w_j = v_j$ for j odd and $w_j \in V_j$ for j even, for all $j \in I$, such that the color sequence $\phi(w_a), \ldots, \phi(w_{a+4k+1})$ is a repetition (of size $2k + 1$). Since this cannot happen, it follows that there exists an index $i \in [a, a + 2k]$ for which one of the above two conditions is not satisfied. For every such index i, we say that the pair (p, q) is a *witness* (for interval I), where $\{p, q\} = \{i, i + 2k + 1\}$ with p odd and q even.

Next, consider an even index $q \in [2n]$. Observe that $|[\ell n] - \phi(V_q)| \leqslant \ell - 1$, since $\phi(V_q)$ contains at least one color from each ℓ-subset of $[\ell n]$. Combining this with the fact that vertices v_p with odd index $p \in [2n]$ have pairwise disjoint lists, we deduce that there are at most $\ell - 1$ odd indices $p \in [2n]$ such that the pair (p, q) is a witness. Summing up over every even index $q \in [2n]$, it follows that there are at most

$$n(\ell - 1)$$

distinct witnesses in total.

Now consider a witness (p, q) and let $|p - q| = 2k + 1$. The pair (p, q) is a witness for at most $2k + 1$ intervals $I \subseteq [2n]$ of the form $I = [a, a + 4k + 1]$. Since there are exactly $2n - 4k - 1$ intervals I of the latter form and each interval of that form must have a witness, it follows that the number of witnesses (p, q) with $|p - q| = 2k + 1$ is at least $\frac{2n-4k-1}{2k+1}$. Summing up over every possible value of k (that is, $k = 0, 1, \ldots, \lfloor (n - 1)/2 \rfloor$),

we obtain that the total number of witnesses is at least

$$\begin{aligned}\sum_{k=0}^{\lfloor (n-1)/2 \rfloor} \frac{2n-4k-1}{2k+1} &\geqslant \sum_{k=0}^{\lfloor (n-1)/2 \rfloor} \left(\frac{n}{k+1} - 2 \right) \\ &\geqslant n \left(\sum_{k=1}^{\lfloor (n+1)/2 \rfloor} \frac{1}{k} \right) - 2 \left\lfloor \frac{n+1}{2} \right\rfloor \\ &\geqslant n \ln \left(\left\lfloor \frac{n+1}{2} \right\rfloor \right) - (n+1) \\ &\geqslant n \ln \left(\frac{n}{2} \right) - (n+1) \\ &\geqslant n \ln n - 3n.\end{aligned}$$

It follows that

$$n(\ell - 1) \geqslant n \ln n - 3n,$$

which contradicts the assumption that $n > e^{\ell+2}$. □

4 Trees of bounded pathwidth

A *path-partition* of a tree T is a pair $(\mathcal{T}, \mathcal{P})$ where $\mathcal{T}$ is a rooted tree and $\mathcal{P}$ is a collection $\{P_x : x \in V(\mathcal{T})\}$ of vertex-disjoint paths of T which collectively partition the vertex set of T and such that $xy \in E(\mathcal{T})$ if and only if there is an edge between a vertex from P_x and a vertex from P_y in T. Observe that a consequence of this definition is that $\mathcal{T}$ is a minor of T. The *root-path* of $(\mathcal{T}, \mathcal{P})$ is the path P_x where x is the root of $\mathcal{T}$. Now, consider a path P_x with x distinct from the root. The path P_x has a *center*, defined as the endpoint in P_x of the edge in T linking P_x to P_y where y is the parent of x in $\mathcal{T}$. The *height* of the path-partition $(\mathcal{T}, \mathcal{P})$ is the height of $\mathcal{T}$.

When considering a path-partition $(\mathcal{T}, \mathcal{P})$ of a tree T, it will be useful to embed T itself in the plane in a way that is 'faithful' to the path-partition. This leads to the following definition: An embedding of T in the plane is *faithful* to the path-partition $(\mathcal{T}, \mathcal{P})$ if each path in $\mathcal{P}$ is drawn horizontally, and contracting each such path into one of its vertices we obtain some plane embedding of $\mathcal{T}$, with its root drawn at the bottom and its edges going up. See Figure 1 for an illustration. As the paths in $\mathcal{P}$ are drawn horizontally, they have a natural orientation from left to right. Every edge e of T is either *horizontal* or *vertical*, depending on whether e belongs to some path in $\mathcal{P}$ or not.

Our motivation for considering path-partitions is the following lemma.

Lemma 4. *Every tree of pathwidth k has a path-partition of height at most $2k$.*

Proof. We prove the following stronger statement: For every tree T of pathwidth k and every vertex $u \in V(T)$, there is a path-partition of T of height at most $2k$ with u in the root-path.

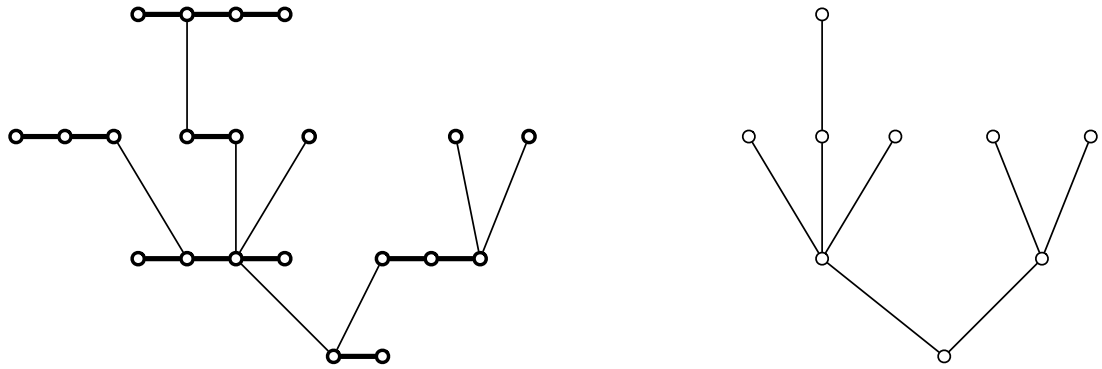

Figure 1: Left: A tree T faithfully embedded according to some path-partition $(\mathcal{T}, \mathcal{P})$ (the paths in $\mathcal{P}$ are drawn in bold). Right: The rooted tree $\mathcal{T}$.

The proof is by induction on k. For $k = 0$, the tree T consists of the single vertex u. Clearly, it has a path-partition of height 0 with u in the root-path. Now suppose $k > 0$ for the inductive case. Let $(P, \mathcal{C})$ be a path decomposition of T of width k, where $T_v \in \mathcal{C}$ denotes the path associated to vertex $v \in V(T)$. Enumerate the vertices of the path P indexing the path decomposition as $p_1, \ldots, p_n$, in order. We may assume without loss of generality that there are (non-necessarily distinct) vertices $x, y \in V(T)$ such that $p_1 \in V(T_x)$ and $p_n \in V(T_y)$ (otherwise P could be shortened). Let $Q_1 = v_1 v_2 \ldots v_m$ denote the unique path in T between $x = v_1$ and $u = v_m$ in T. Let z be the vertex of Q_1 that is closest to y in T. Let $Q_2 = w_1 w_2 \ldots w_{m'}$ denote the unique path in T between $y = w_1$ and $z = w_{m'}$ in T. Notice that $V(T_{v_i}) \cap V(T_{v_{i+1}}) \neq \emptyset$ for each $i \in \{1, \ldots, m-1\}$ and that each $V(T_{v_i})$ induces a subpath of P. Similarly, $V(T_{w_i}) \cap V(T_{w_{i+1}}) \neq \emptyset$ for each $i \in \{1, \ldots, m'-1\}$ and each $V(T_{w_i})$ induces a subpath of P. Thus we deduce that $\bigcup_{i=1}^{m} V(T_{v_i}) \cup \bigcup_{i=1}^{m'} V(T_{w_i}) = \{p_1, \ldots, p_n\}$.

Consider the forest $T - (V(Q_1) \cup V(Q_2))$ and let $D_1, \ldots, D_c$ denote its components. Observe that each tree D_j ($j \in \{1, \ldots, c\}$) has pathwidth at most $k-1$. Indeed, $(P, \{T_v \mid v \in V(D_j)\})$ is a path decomposition of D_j, and for each $i \in \{1, \ldots, n\}$ we have $|\{v \in V(D_j) \mid p_i \in V(T_v)\}| \leqslant |\{v \in V(T) \mid p_i \in V(T_v)\}| - 1 \leqslant k-1$.

For each $j \in \{1, \ldots, c\}$, let d_j denote the unique vertex of D_j having a neighbor in $Q_1 \cup Q_2$ in T. By induction, each tree D_j ($j \in \{1, \ldots, c\}$) has a path-partition $(\mathcal{T}_j, \mathcal{P}_j)$ of height at most $2(k-1)$ such that d_j in the root-path. Let r_j denote the root of $\mathcal{T}_j$.

We construct a path-partition $(\mathcal{T}, \mathcal{P})$ of T as follows: $\mathcal{T}$ consists of the disjoint union of $\mathcal{T}_1, \ldots, \mathcal{T}_c$ plus two extra vertices q_1 and q_2 with q_1 the root of $\mathcal{T}$ and q_2 a child of q_1. The paths associated to q_1 and q_2 are Q_1 and Q_2, respectively. For each $j \in \{1, \ldots, c\}$, we make q_1 or q_2 adjacent to r_j, depending whether d_j has a neighbor in Q_1 or Q_2 in T. It is easy to verify that $(\mathcal{T}, \mathcal{P})$ is a path-partition of T of height at most $2(k-1)+2 = 2k$. □

The fact that trees of bounded pathwidth have path-partitions of bounded height is a natural observation. It is thus likely that this observation was made before though we are not aware of any relevant reference. We also note that we made no effort to optimize

the bound in Lemma 4 and we do not know whether the factor 2 is unavoidable.

Let T be a tree and fix a path-partition $(\mathcal{T}, \mathcal{P})$ of T. Suppose further that T is embedded in the plane faithfully to $(\mathcal{T}, \mathcal{P})$. We use the following terminology when discussing paths in T. First, every path $P \in \mathcal{P}$ has a corresponding *level*, which is defined as the height of the corresponding vertex in $\mathcal{T}$. By extension, every vertex of T has a level, the level of the path in $\mathcal{P}$ it belongs to. Define the *base* of an arbitrary path P in T as the subpath induced by the vertices of P of minimum level. Since the base of P is a subpath of a path in $\mathcal{P}$, its vertices are ordered from left to right by the plane embedding of T. The path P is said to be *ascending* if at least one of its two endpoints belongs to its base. (We note that in particular all paths in $\mathcal{P}$ are ascending, even though they are drawn horizontally in the embedding of T.) Each ascending path P in T has a *source*, defined as the endpoint of P that is in the base of P; in case both endpoints are in the base, the left-most one is selected as the source. We typically think of ascending paths P as being directed from their source to their other endpoint so that the notion of ith vertex of P is well defined, the first vertex being the source. An ascending path P with at least two vertices either *goes right* or *goes left* or *goes up*, depending on whether the second vertex of P is on the base and to the right of the source, or on the base and to the left of the source, or is one level higher.

Next we generalize the notion of repetition as follows. A *near repetition* is a sequence of the form $x_1 \dots x_r y_1 \dots y_g x_1 \dots x_r$, where $r \geqslant 1$ is its *length*, and $g \geqslant 0$ is said to be its *gap*. (Thus for $g = 0$ this is the usual notion of repetition.) Now, let us return to our tree T from the previous paragraph, and let ϕ denote an arbitrary coloring of its vertices. A slightly technical but key definition for our purposes is the following: An ascending path P of T is said to be *ϕ-bad* if, enumerating its vertices as $v_1, v_2, \dots, v_p$ starting from its source, the sequence $\phi(v_1)\phi(v_2)\dots\phi(v_p)$ forms a near repetition $x_1 \dots x_r y_1 \dots y_g x_1 \dots x_r$ of length r and gap g where at most r vertices from $v_{r+1}, \dots, v_{r+g}$ lie in the base of P. (That is, either $g \leqslant r$, or $g > r$ but at most r vertices from the 'gap' section are in the base of P.) An ascending path that is not ϕ-bad is said to be *ϕ-good*. We sometimes drop ϕ when using these two adjectives if the coloring ϕ they refer to is clear from the context

Equipped with these definitions we may now state the following technical lemma, which turns out to be the heart of the proof.

Lemma 5. *There is a function $f : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ such that, for every $\ell \geqslant 1$, every $h \geqslant 0$, and every tree T faithfully embedded according to a path-partition $(\mathcal{T}, \mathcal{P})$ of T of height h with lists L_v $(v \in V(T))$ of colors of size $f(\ell, h)$, one can find sublists $S_v \subseteq L_v$ $(v \in V(T))$ of size ℓ such that, for every coloring ϕ of T with these sublists, every ascending path of T is ϕ-good.*

In order to motivate Lemma 5, we show that with only a little extra effort (greedy coloring from the sublists) it implies Theorem 2.

Proof of Theorem 2 (assuming Lemma 5). Let T be a tree of pathwidth k and let L_v $(v \in V(T))$ be a list assignment for the vertices of T where each list has size $b(k) := f(2k+1, 2k)$, where f is the function from Lemma 5. By Lemma 4 there is a path-partition $(\mathcal{T}, \mathcal{P})$ of T of height at most $2k$. By Lemma 5 there are sublists $\{S_v\}_{v \in V(T)}$

with $S_v \subseteq L_v$ and $|S_v| = 2k+1$ for each vertex $v \in V(T)$, such that in any coloring ϕ of T with these sublists, all ascending paths in T are ϕ-good.

We define a nonrepetitive coloring ϕ of T with the lists S_v ($v \in V(T)$) in a greedy manner. We color the vertices of T one by one in non-decreasing order of their levels. Let $v \in V(T)$ be a vertex under consideration. Let $v_1 v_2 \dots v_p$ denote the shortest path in T from $v_1 = v$ to the root-path (thus it enters the root-path in vertex v_p). Recall that every edge of the form $v_{i-1}v_i$ with $i \in \{2, \dots, p\}$ is either horizontal or vertical in T. Let

$$G(v) := \big\{ v_i \mid i \in \{2, \dots, p\} \text{ and } v_{i-1}v_i \text{ is a vertical edge} \big\}.$$

Each vertex in $G(v)$ is said to be a *guard* for vertex v. Note that $|G(v)|$ is exactly the level of vertex v in T, and thus in particular $|G(v)| \leqslant 2k$. We color v as follows: Let $\phi(v)$ be an arbitrarily chosen color from the non-empty set $S(v) - \phi(G(v))$. (Here, $\phi(G(v))$ denotes the set of colors used for vertices in $G(v)$; note that these vertices are already colored since they lie on lower levels.)

We claim that ϕ is a nonrepetitive coloring of T. Arguing by contradiction, suppose that there is a repetitively colored path $P = v_1 \dots v_p w_1 \dots w_p$. Consider the edge $e = v_p w_1$. First we show that e belongs to the base of P. Suppose not, and consider the shortest subpath of P that includes the edge e and has one endpoint in the base of P. Reversing P if necessary, we may assume without loss of generality that this subpath is of the form $v_p w_1 \dots w_m$, with w_m being the only vertex on the base. Observe that $w_{m-1}w_m$ is a vertical edge. This implies that w_m is a guard for all the vertices in $\{v_1, \dots, v_p, w_1, \dots, w_{m-1}\}$. In particular, the color $\phi(w_m)$ cannot have been used for vertex v_m since $\phi(v_m) \in S(v_m) - \phi(G(v_m))$, contradicting the fact that $\phi(w_m) = \phi(v_m)$. Therefore, the edge e must lie in the base of P.

Let ℓ and r be the number of vertices in $\{v_1, \dots, v_p\}$ and $\{w_1, \dots, w_p\}$, respectively, that are in the base of P. Reversing P if necessary, we may assume without loss of generality that $\ell \leqslant r$. Consider the path $P' = w_r w_{r-1} \dots w_1 v_p \dots v_1$. Observe that P' is an ascending path as one of its endpoints, namely w_r, is in the base of P'. Now, $\phi(w_r)\phi(w_{r-1}) \dots \phi(w_1)\phi(v_p) \dots \phi(v_1)$ is a near repetition of length r with gap $p - r$, and exactly ℓ vertices from the gap section are in the base of P'. Since $\ell \leqslant r$, we deduce that P' is ϕ-bad, contradicting the fact that every ascending path is ϕ-good. □

An *arborescence* is a rooted directed tree where the edges are directed away from the root. It will be convenient to consider arborescences that are embedded in the plane without edge crossings in such a way that the root is drawn at the bottom and all arcs go up (thus the source of an arc is drawn below its sink), which we simply call *plane* arborescences. The *height* of a vertex in an arborescence is defined as its distance to the root, thus in particular the root has height 0. The *rightmost path* of a plane arborescence is the path obtained by starting from the root and always taking the rightmost arc going up, until reaching a leaf.

We classify directed paths in a plane arborescence A as being good or bad w.r.t. a given coloring ϕ of A, similarly as for ascending paths: Say that a directed path P is *ϕ-bad* if, enumerating its vertices as $v_1 v_2 \dots v_p$ in order, the sequence $\phi(v_1)\phi(v_2) \dots \phi(v_p)$

can be written as a near repetition $x_1 \dots x_r y_1 \dots y_g x_1 \dots x_r$ of length r and gap g where at most r vertices from $v_{r+1}, \dots, v_{r+g}$ lie on the rightmost path of A. (That is, either $g \leqslant r$, or $g > r$ but at most r vertices from the 'gap' section are on the rightmost path.) If the directed path P is not ϕ-bad then it is *ϕ-good*.

Lemma 6. *Let $\ell \geqslant 1$, let A be a plane arborescence, and let L_v ($v \in V(A)$) be lists of colors of size $32\ell^3 + 1$. Then one can find sublists $S_v \subseteq L_v$ ($v \in V(A)$) of size ℓ such that, for every coloring ϕ of A with these sublists, every directed path starting on the rightmost path is ϕ-good.*

The interest of Lemma 6 is that Lemma 5 can be proved by iterated applications of Lemma 6, as we now show.

Proof of Lemma 5 (assuming Lemma 6). The function $f(\ell, h)$ that will be used is defined inductively on h as follows: $f(\ell, 0) := 32\ell^3 + 1$, and $f(\ell, h) := f(32(32\ell^3 + 1)^3 + 1, h - 1)$ for $h > 0$.

Let T be a tree with a path-partition $(\mathcal{T}, \mathcal{P})$ of height h and let L_v ($v \in V(T)$) be lists of colors of size $f(\ell, h)$. Suppose further that T is faithfully embedded according to $(\mathcal{T}, \mathcal{P})$. We prove the lemma by induction on h. For the base case of the induction, $h = 0$, we observe that T is then a path and all ascending paths in T are simply subpaths of T. As $f(\ell, 0) = 32\ell^3 + 1$, by Lemma 6 there are sublists $S_v \subseteq L_v$ for each vertex $v \in V(T)$ with $|S_v| = \ell$ such that, for every coloring ϕ of T with these sublists, all ascending paths of T are ϕ-good, as required. (As expected, when applying Lemma 6 we first turn the path T into an arborescence by directing it from left to right.)

For the inductive case $h > 0$, let x be the root of $\mathcal{T}$ and let P_x be the root-path of T. Let also $D_1, \dots, D_c$ be the components of the forest $T - V(P_x)$. (Note that there is at least one component.) For each $i \in \{1, \dots, c\}$, the path-partition $(\mathcal{T}, \mathcal{P})$ induces in a natural way a path-partition $(\mathcal{T}_i, \mathcal{P}_i)$ of D_i of height at most $h - 1$, with $\mathcal{T}_i$ rooted at the only vertex that is a neighbor of x in $\mathcal{T}$. Since $f(\ell, h) = f(32(32\ell^3 + 1)^3 + 1, h - 1)$, applying induction on D_i we obtain for each vertex $v \in V(D_i)$ a sublist $S'_v \subseteq L_v$ of size $32(32\ell^3 + 1)^3 + 1$ such that, for every coloring ϕ of D_i with these sublists, every ascending path of D_i is ϕ-good.

Next, for each vertex $v \in V(P_x)$ let S'_v be an arbitrary subset of L_v of size $32(32\ell^3 + 1)^3 + 1$. Thus, every vertex v of T now has a corresponding sublist $S'_v \subseteq L_v$ of size $32(32\ell^3 + 1)^3 + 1$. Moreover, given any coloring ϕ of the tree T with these sublists, the only ascending paths that could possibly be ϕ-bad are those having their sources in P_x. We shall refer to these ascending paths as the *risky paths* of T.

Enumerate the vertices of the root-path P_x as $v_1 v_2 \dots v_n$, from left to right. Define two plane arborescences A and A' from T by rooting T at v_1 and v_n, respectively, and ensuring that P_x is a prefix of the rightmost path in both instances. Note that the rightmost path of A could extend beyond P_x (in case v_n is not a leaf of T), and the same is true for the rightmost path of A' (if v_1 is not a leaf). What is important for our purposes is to observe that each risky path of T starts on the rightmost path in both A and A'. Observe also that each risky path of T that goes right (left) is a directed path in A (respectively A'), and risky paths that go up are directed in both A and A'.

First, apply Lemma 6 on A with list assignment S'_v ($v \in V(A)$), giving for each vertex $v \in V(T)$ a sublist $S''_v \subseteq S'_v \subseteq L_v$ of size $32\ell^3 + 1$. Next, apply Lemma 6 on A' with list assignment S''_v ($v \in V(A)$), giving for each vertex $v \in V(T)$ a sublist $S_v \subseteq S''_v \subseteq S'_v \subseteq L_v$ of size ℓ. Since every risky path of T is mapped to a directed path starting on the rightmost path in A or A', by the properties of the sublists S''_v and S_v ($v \in V(T)$) guaranteed by Lemma 6 we know that, for every coloring ϕ of T with the lists S_v ($v \in V(T)$), all risky paths of T are ϕ-good. Therefore, the lists S_v ($v \in V(T)$) have the desired properties. □

It remains to prove Lemma 6. As alluded to in the introduction, we will do so by adapting the algorithmic method used in [6, 10, 13].

Proof of Lemma 6. Let $N := 32\ell^3 + 1$ denote the size of the lists. For $v \in V(A)$, let $\mathrm{UP}(v)$ denote the set of vertices $w \in V(A)$ that can be reached via a directed path from v in A. (Note that $v \in \mathrm{UP}(v)$.) In the proof, we will often abbreviate 'subset of size k' and 'sublist of size k' into *'k-subset'* and *'k-sublist'*, respectively.

We define a simple randomized algorithm, Algorithm 1, that tries to find an ℓ-sublist S_v of L_v for each vertex $v \in V(A)$ such that, for every coloring ϕ of A with these sublists, every directed path starting on the rightmost path of A is ϕ-good. The following informal description of the algorithm is complemented by the more formal description given in Algorithm 1. The algorithm explores the arborescence A via a depth-first, left-to-right search starting from the root. The algorithm maintains at all time ℓ-sublists $S_v \subseteq L_v$ for all vertices v encountered *before* the current vertex u in the depth-first search of A. These sublists have the following property: For every coloring ϕ of these vertices with these sublists (ϕ being thus a partial coloring of A), every directed path starting on the rightmost path of A that is fully colored is ϕ-good. We say that such a partial sublist assignment is *valid*.

Next, the algorithm treats the current vertex u and tries to maintain the above property. To do so, the algorithm first chooses an ℓ-sublist $S_u \subseteq L_u$ uniformly at random. If this new sublist S_u triggers the existence of a ϕ-bad path in A for some (partial) coloring ϕ with the current sublists—that is, the current sublist assignment is no longer valid—it erases some of these sublists as follows: Say $v_1 \dots v_{2r+g}$ with $v_{2r+g} = u$ is a ϕ-bad path with color sequence $\phi(v_1) \dots \phi(v_{2r+g})$ of the form $x_1 \dots x_r y_1 \dots y_g x_1 \dots x_r$. The algorithm then erases the choice for the list S_v for all vertices v contributing to the second occurrence of the repeated sequence and their descendants, that is, for all $v \in \mathrm{UP}(v_{r+g+1})$. At the next iteration, v_{r+g+1} becomes the new current vertex, that is, the next vertex to be treated. Notice that this makes the algorithm backtrack a number of steps w.r.t. the depth-first left-to-right search of A.

If on the other hand, the new sublist S_u does not trigger any such bad configuration, then the current sublist assignment remains valid. In this case, before proceeding to the next random choice the algorithm first tries to extend the current sublist assignment deterministically as much as possible. (While it might not be clear at first glance why this deterministic extension step is needed, we remark that it is actually a key feature of the algorithm without which we could not do the analysis below.) This is done as follows: The algorithm considers the children $u_1, \dots, u_k$ of u one by one in left-to-right

order, until a *problematic* child is identified: When considering u_j, the algorithm checks whether there exist ℓ-subsets $S_v \subseteq L_v$ for all $v \in \mathrm{UP}(u_j)$ such that, taken together, they extend the current sublist assignment in such a way that it remains valid. If these subsets exist, the current sublist assignment is extended in this way to the whole subtree rooted at u_j, and the algorithm considers the next child of u. (If there are more than one valid choice for these sublists, the algorithm chooses one according to a deterministic rule.) If no such extension of the current sublist assignment can be found for vertices in $\mathrm{UP}(u_j)$, then u_j is identified as being a problematic child of u, and u_j becomes the next vertex to be treated. Observe that this effectively makes the algorithm proceed with the depth-first left-to-right search of A for some number of steps.

Let us make some observations concerning the algorithm: Right at the beginning, after selecting a sublist for the root of A, two situations can occur: (1) No child of the root is problematic. Thus a valid sublist assignment for all vertices of A has been found, and the lemma is proved. (2) Some child of the root is problematic. In this case, it is important to observe that later on each vertex u for which a random sublist S_u is chosen was problematic when its parent was considered. It follows in turn that some child of u will be problematic, since otherwise we could extend the sublist assignment to the whole subtree rooted at u.

To summarize, we may assume that we are in case (2) at the beginning, since otherwise we are done. This implies that the vertex u that is currently being treated by the algorithm always has a problematic child. Moreover, the algorithm *will never stop*, simply because while it can erase the choices of sublists for some vertices of A it cannot do so for the root, as is easily checked. Our proof will then proceed in the following way: We run the algorithm until it made M random choices of sublists and then stop it, where M will be some large number which is a function of $|V(A)|$ and ℓ. We then carefully set up a concise description (called *log*) of its execution that is precise enough to allow us to recover from it all random choices that were made by the algorithm. Finally, we count the number of distinct logs that can occur after M random choices, and show that, for sufficiently large M, this number is strictly less than $\binom{N}{\ell}^M$. From this we deduce that not all sequences of M random choices of sublists can occur in case (2). In other words, there is a choice for the sublist of the root of A leaving us in case (1), which then finishes the proof.

This concludes our informal description of the algorithm, see Algorithm 1 for the pseudo code. A few remarks about the latter are in order: First, we assume that the ℓ-subsets of L_v have been enumerated for each $v \in V(A)$, so that the j-th ℓ-subset of L_v is well defined for $j \in \left[\binom{N}{\ell}\right]$. This ordering also induces an ordering on every subcollection of the collection of ℓ-subsets of L_v. We also use this enumeration in the proof. Second, for simplicity we model the random choices made by the algorithm by a sequence $r_1, r_2, \ldots, r_M$ of numbers given in input, each between 1 and $\binom{N}{\ell}$, where r_i will be the number used for the i-th random choice. We call this sequence the *random input*. Third, in line 8, the ϕ-bad path is chosen according to some fixed rule. Similarly, in line 15, the sublists S_v are chosen according to some fixed rule. (In each case, the actual rule is irrelevant, as long as it is deterministic.)

In the following, by the i-th iteration of the algorithm, we mean the i-th iteration

Algorithm 1: Attempts to find sublists S_v of L_v for all $v \in V(A)$, each of size ℓ, such that for every coloring ϕ of A with these sublists, every directed path starting on the rightmost path of A is ϕ-good.

```
1   input: Lists L_v for all v ∈ V(A) and random input r_1, r_2, ..., r_M
2   i ← 1
3   u ← root of A
4   S_v ← undefined for each v ∈ V(A)
5   while i ⩽ M do
6       S_u ← r_i-th subset of size ℓ of L_u
7       if there is a φ-bad path starting on the rightmost path of A for some coloring
        φ with the lists S_v (v ∈ V(A)) then
8           let v_1 ... v_{2r+g} with v_{2r+g} = u be a φ-bad path and let
              φ(v_1) ... φ(v_{2r+g}) be a sequence of the form x_1 ... x_r y_1 ... y_g x_1 ... x_r
9           S_v ← undefined for each v ∈ UP(v_{r+g+1})
10          u ← v_{r+g+1}
11      else
12          let u_1, ..., u_k denote the children of u ordered from left to right
13          j ← 1
14          while  j ⩽ k and there is a valid extension of the current sublist
            assignment by some S_v ⊆ L_v for all v ∈ UP(u_j) do
15              choose such sublists S_v ⊆ L_v for all v ∈ UP(u_j)
16              j ← j + 1
17          if j = k + 1 then
                (in this case i = 1 and u is the root of A)
18              return S_v for all v ∈ V(A)
19          else u ← u_j
20      i ← i + 1
21  report failure
```

of the `while` loop. We call operations in lines 8-10 a *retraction* of the near repetition $x_1 \dots x_r y_1 \dots y_g x_1 \dots x_r$. With a slight abuse of terminology, we will also say that the corresponding ϕ-bad path has been *retracted*.

From now on we argue by contradiction and suppose that the desired sublists for the vertices of A do not exist. In other words, we assume that every choice for the sublist of the root at the beginning of the algorithm leaves us in case (2) described above. In particular, for all M and all random inputs $r_1, r_2, \dots, r_M$, Algorithm 1 runs for M steps and then reports failure.

We are going to create a concise description of what Algorithm 1 does during the M steps of its execution. This description is completely determined by the lists and the random input. We see the lists L_v ($v \in V(A)$) as being fixed and thus treat the description as a function of the random input $r_1, r_2, \dots, r_M$. The description, which we call an *M-log*,

consists of a 4-tuple $(D, \mathcal{S}, B, \Gamma)$ defined as follows:

(i) $D = (d_1, \ldots, d_M)$ and d_i ($i \in [M]$) is the height of the vertex u in Algorithm 1 at the *end* of iteration i, when reaching line 20 after u was updated in the `if`–`else` block. (Thus, for $i < M$, d_i is simply the height of vertex u at the *beginning* of iteration $i + 1$.)

(ii) $\mathcal{S} : V(A) \to \left[\binom{N}{\ell}\right] \cup \{\texttt{undefined}\}$ is a function encoding the final partial sublist assignment to vertices of A at the end of iteration M. More precisely, for each $v \in V(A)$ we have $\mathcal{S}(v) = j$ if sublist S_v is defined at that moment and S_v is the j-th ℓ-subset of L_v, and $\mathcal{S}(v) = \texttt{undefined}$ if S_v is not defined.

(iii) $B = (b_1, \ldots, b_M)$ and b_i ($i \in [M]$) are defined as follows: If no bad path was retracted during the i-th iteration then $b_i = 0$. Otherwise, b_i is the number of vertices in the prefix $v_1 \ldots v_{r+g}$ of the retracted bad path $v_1 \ldots v_{2r+g}$ that are on the rightmost path of A. Observe that $b_i \leqslant 2r$ in the latter case since at most r vertices from the gap section of a bad path lie on the rightmost path of A.

(iv) $\Gamma = (\Gamma_1, \ldots, \Gamma_M)$, where Γ_i ($i \in [M]$) is defined as follows. If no bad path was retracted during the i-th iteration then $\Gamma_i = \texttt{undefined}$. Otherwise, letting $v_1 \ldots v_{2r+g}$ be the retracted bad path, we set $\Gamma_i = (\gamma_1, \ldots, \gamma_r)$ with γ_j ($j \in [r]$) defined as follows: Let $\mathcal{X}_j$ denote the collection of ℓ-subsets of $L_{v_{r+g+j}}$ that have a *non-empty intersection with* S_{v_j}. Then γ_j is the index of the set $S_{v_{r+g+j}}$ in the collection $\mathcal{X}_j$.

Now, our aim is to bound from above the number of distinct M-logs $(D, \mathcal{S}, B, \Gamma)$ by a relatively small function of $M, |V(A)|$, and ℓ. Recall that the lists L_v ($v \in V(A)$) are fixed, thus ℓ and $|V(A)|$ are fixed, and only M and the random input $r_1, \ldots, r_M$ vary. There are exactly $\binom{N}{\ell}^M$ distinct random inputs of length M, and our goal in the following analysis is to deduce that there are $o\left(\binom{N}{\ell}^M\right)$ distinct M-logs. (The asymptotic notation is to be interpreted with respect to the variable M of course.) This is then a contradiction for M large enough, as mentioned earlier.

We start by estimating the number of M-tuples $D = (d_1, \ldots, d_M)$. Each sequence $D = (d_1, \ldots, d_M)$ can be injectively mapped to its sequence of differences $(d_2 - d_1, \ldots, d_M - d_{M-1})$. (Note that $d_1 = 1$.) All numbers in this new sequence belong to the set $\{1, 0, -1, -2, \ldots\}$, as is easily seen. Next we transform that sequence into yet another sequence by replacing each number k by 1 followed by $1-k$ consecutive occurrences of -1. For instance, the sequence of differences $(1, 1, 1, 1, 1, -2, -1, 1)$ gets mapped to $(1, 1, 1, 1, 1, 1, -1, -1, -1, 1, -1, -1, 1)$. It is easy to see that the second transformation is also injective. The resulting sequence D' is a sequence over the alphabet $\{-1, 1\}$. The number of 1's in D' corresponds to the number of times the algorithm assigns a value to some variable S_u in line 6, and is thus equal to the number of iterations, that is, M. The number of -1's in D' is the sum of all values of r over all bad paths $v_1 \ldots v_{2r+g}$ considered in lines 8-9 during the execution. One can see this as the number of times the algorithm

'erases' a value of S_v for some $v \in V(A)$ that was set earlier using the random input (note that an execution of line 9 erases r such values). Thus, this number is at most the total number of executions of line 6, that is, the number of 1's in D', which is M. Hence, D' has size between M and $2M$, and there are at most $2^M + 2^{M+1} + \cdots + 2^{2M} \leqslant 2^{2M+1}$ such sequences D'.

Next we bound the number of different functions $\mathcal{S}$. Note that this number depends only on N, ℓ, and $|V(A)|$, so it can be treated as a constant w.r.t. M. We denote this number by c (its exact value being irrelevant for the analysis).

Now we turn our attention to the number of possible tuples $B = (b_1, \ldots, b_M)$ in an M-log $(D, \mathcal{S}, B, \Gamma)$. Recall that if $b_i \neq 0$ then $b_i \leqslant 2r$, where r is the size of the repeated part in the near repetition retracted during the i-th iteration. Hence, the sum of the b_i's is at most twice the total number of -1's in D'. This implies that $b = \sum_{i=1}^{M} b_i \leqslant 2M$. The number of such sequences is easily seen to be at most 2^{3M}, as one can encode each such sequence as a binary word $1^{b_1}01^{b_2}01^{b_3}0 \ldots 1^{b_M}01^{2M-b}$ of length $3M$ (the i-th section $1^{b_i}0$ encodes the number b_i for $i = 1, \ldots, M$, and the padding 1^{2M-b} at the end of the word ensures that the length is $3M$).

It remains to estimate the number of possible tuples $\Gamma = (\Gamma_1, \ldots, \Gamma_M)$ which can occur for fixed D, $\mathcal{S}$, and B. Consider an index $i \in [M]$. If $\Gamma_i \neq \texttt{undefined}$, then each number in the sequence Γ_i is the index of some ℓ-subset of some list L_v ($v \in V(A)$) among those that intersect some other list S of size ℓ. Clearly, the number of ℓ-subsets of L_v intersecting S is at most $\binom{\ell}{1}\cdot\binom{N-1}{\ell-1} + \binom{\ell}{2}\cdot\binom{N-2}{\ell-2} + \cdots + \binom{\ell}{\ell-1}\cdot\binom{N-(\ell-1)}{1} + \binom{\ell}{\ell}\cdot\binom{N-\ell}{0}$. Since N is much bigger than ℓ (recall that $N > 32\ell^3$), each term in the previous sum is upper bounded by the first term $\ell \cdot \binom{N-1}{\ell-1}$, and we obtain that the sum is at most $\ell^2 \cdot \binom{N-1}{\ell-1}$. (Of course, this is a rather crude upper bound but it is good enough for our purposes.) Hence all numbers in the sequence Γ_i are between 1 and $\ell^2 \cdot \binom{N-1}{\ell-1}$. Note also that the length of the sequence Γ_i is exactly the length of the near repetition retracted during the i-th iteration. Hence, given D we know exactly which Γ_i are defined and what are their lengths. The sum of these lengths is the total number of -1's in D', which is at most M. Therefore, for a fixed D there can be at most $\left(\ell^2 \cdot \binom{N-1}{\ell-1}\right)^M$ distinct sequences Γ.

Putting all the previous observations together, we deduce that the number of distinct tuples $(D, \mathcal{S}, B, \Gamma)$ is at most

$$\begin{aligned} 2^{2M+1} \cdot c \cdot 2^{3M} \cdot \left(\ell^2 \binom{N-1}{\ell-1}\right)^M &= O\left(32^M \cdot \left(\frac{\ell^3}{N}\binom{N}{\ell}\right)^M\right) \\ &= o\left(\binom{N}{\ell}^M\right), \end{aligned}$$

as desired. (The $o(\cdot)$ follows from the fact that $N > 32\ell^3$.) This shows that, if M is sufficiently large, then the number of possible M-logs is strictly smaller than $\binom{N}{\ell}^M$, the number of random inputs of length M. To obtain the desired contradiction, it remains

to show that runs of the algorithm on different sources produce distinct M-logs, that is, that any M-log $(D, \mathcal{S}, B, \Gamma)$ uniquely determines the random input used by the algorithm to produce it. This is exactly what we show next.

Consider an M-log $(D, \mathcal{S}, B, \Gamma)$ and let $r_1, \ldots, r_M$ be *any* random input that can lead to its production. We prove that $r_1, \ldots, r_M$ are uniquely determined by induction on M. This is clearly true if $M = 1$, since the function $\mathcal{S}$ tells us explicitly which sublist was chosen for the root of A. So assume $M > 1$ for the inductive case. Let $D = (d_1, \ldots, d_M)$, $B = (b_1, \ldots, b_M)$, and $\Gamma = (\Gamma_1, \ldots, \Gamma_M)$.

First suppose that $d_M = d_{M-1} + 1$. Then no near repetition was retracted during the M-th iteration. The vertex u that was the current vertex at the beginning of the M-th iteration is determined by the function $\mathcal{S}$: It is the *last* vertex $w \in V(A)$ in the depth-first left-to-right search order from the root such that $\mathcal{S}(w) \neq \texttt{undefined}$ that has a child w' with $\mathcal{S}(w') = \texttt{undefined}$. (Note that the first such child w' is the problematic child u_j of u identified when exiting the inner while-loop.) Now, observe that r_M is simply the index of $\mathcal{S}(u)$ among ℓ-subsets of L_u, and is thus completely determined by our M-log $(D, \mathcal{S}, B, \Gamma)$.

Having determined r_M, we can use the inductive hypothesis to deduce that $r_1, \ldots, r_{M-1}$ are also fully determined by the log as follows. Let $(D^*, \mathcal{S}^*, B^*, \Gamma^*)$ be the $(M-1)$-log resulting from the execution of the algorithm for $M - 1$ iterations on random input $r_1, \ldots, r_{M-1}$. By induction, the latter sequence is uniquely determined by the $(M-1)$-log $(D^*, \mathcal{S}^*, B^*, \Gamma^*)$. Hence, it is enough to show that $(D^*, \mathcal{S}^*, B^*, \Gamma^*)$ is in turn uniquely determined by our initial M-log $(D, \mathcal{S}, B, \Gamma)$. Clearly,

$$\begin{aligned} D^* &= (d_1, \ldots, d_{M-1}), \\ B^* &= (b_1, \ldots, b_{M-1}), \\ \Gamma^* &= (\Gamma_1, \ldots, \Gamma_{M-1}). \end{aligned}$$

As for $\mathcal{S}^*$, it is simply obtained from $\mathcal{S}$ by letting the value of each vertex $v \in \mathrm{UP}(u)$ be `undefined`. That is,

$$\mathcal{S}^*(v) = \begin{cases} \texttt{undefined} & \text{if } v \in \mathrm{UP}(u); \\ \mathcal{S}(v) & \text{otherwise} \end{cases}$$

for each $v \in V(A)$ (recall that $u \in \mathrm{UP}(u)$).

In the case when $d_M = d_{M-1} - r + 1$ with $r > 0$, a near repetition was retracted during the M-th iteration with a repeated part of size r. Here we first show that $r_1, \ldots, r_{M-1}$ are uniquely determined, and then we prove that the same holds for r_M.

Let $(D^*, \mathcal{S}^*, B^*, \Gamma^*)$ denote the $(M-1)$-log resulting from the execution of the algorithm for $M - 1$ iterations on random input $r_1, \ldots, r_{M-1}$. We thus have:

$$\begin{aligned} D^* &= (d_1, \ldots, d_{M-1}), \\ B^* &= (b_1, \ldots, b_{M-1}), \\ \Gamma^* &= (\Gamma_1, \ldots, \Gamma_{M-1}). \end{aligned}$$

Let us show that $\mathcal{S}^*$ is completely determined by $(D, \mathcal{S}, B, \Gamma)$. Let u denote the current vertex at the beginning of the M-th iteration. (Remark: The vertex u can be deduced from the log $(D, \mathcal{S}, B, \Gamma)$, as follows from the discussion below.) During the M-th iteration the sublist S_u of u was assigned the r_M-th ℓ-subset of L_u. This triggered the existence of a bad path $v_1, \ldots, v_{2r+g}$ with $v_{2r+g} = u$, which was subsequently retracted, i.e. S_v was then set to `undefined` for all $v \in \mathrm{UP}(v_{r+g+1})$.

Now, we determine the first $r + g + 1$ vertices $v_1, \ldots, v_{r+g+1}$ of the bad path from the log. First, we show that v_{r+g+1} is easily determined: Since $\mathcal{S}^*$ differs from $\mathcal{S}$ only on the vertices from $\mathrm{UP}(v_{r+g+1})$, vertex v_{r+g+1} is the first vertex v of A in depth-first left-to-right order from the root with $\mathcal{S}(v) = \texttt{undefined}$. (We remark that v_{r+g+1} would have been the next 'current vertex' considered by the algorithm if it were run for an extra iteration.)

Next, we determine the vertices $v_1, \ldots, v_{r+g}$ from the log. Let w be the first vertex on the rightmost path of A that is encountered when walking down towards the root of A from v_{r+g+1} (note that we could have $w = v_{r+g+1}$). By definition of the sequence B, among the $r + g$ vertices $v_1, \ldots, v_{r+g}$ the first b_M vertices are on the rightmost path of A, and none of the remaining $r + g - b_M$ vertices are. Since $b_M \geqslant 1$, we may identify vertex v_1 by starting at vertex w and walking down towards the root of A either b_M steps (if $w = v_{r+g+1}$) or $b_M - 1$ steps (if $w \neq v_{r+g+1}$). Now, since we know vertices v_1, v_{r+g+1} and the integer r, looking at the path from v_1 to v_{r+g+1} gives us all intermediate vertices $v_2, \ldots, v_{r+g}$ (if any), as well as the size g of the gap section. Hence, $v_1, \ldots, v_{r+g+1}$ are determined by the log $(D, \mathcal{S}, B, \Gamma)$, as claimed.

Building on this, we now complete the proof. Let $\Gamma_M = (\gamma_1, \ldots, \gamma_r)$. We first consider a simple case, namely $r = 1$. Then $u = v_{r+g+1}$, and at the beginning of the M-th iteration, the r_M-th ℓ-sublist X of L_u was assigned to variable S_u, which triggered the existence of the bad path. By the definition of the sequence Γ_M, the sublist X is the γ_1-th ℓ-sublist of $L_{v_{r+g+1}}$ having a non-empty intersection with $\mathcal{S}(v_1)$. We can thus deduce X from the log $(D, \mathcal{S}, B, \Gamma)$, and obtain in turn r_M from X. Notice that in this case the sublist assignments to vertices of A at the end of the $(M-1)$-th and at the end of the M-th iterations are exactly the same, that is, $\mathcal{S}^* = \mathcal{S}$. (Indeed, this is what makes the case $r = 1$ simpler.) Hence, $(D^*, \mathcal{S}^*, B^*, \Gamma^*)$ is completely determined by $(D, \mathcal{S}, B, \Gamma)$. By induction, the sequence $r_1, \ldots, r_{M-1}$ is uniquely determined by $(D^*, \mathcal{S}^*, B^*, \Gamma^*)$, and therefore it is also uniquely determined by $(D, \mathcal{S}, B, \Gamma)$. Since we have seen that r_M is uniquely determined as well, this concludes the $r = 1$ case.

Now, suppose that $r \geqslant 2$. We can obtain $\mathcal{S}^*(v_{r+g+1})$ from $(D, \mathcal{S}, B, \Gamma)$, since $\mathcal{S}^*(v_{r+g+1})$ is the γ_1-th ℓ-sublist of $L_{v_{r+g+1}}$ having a non-empty intersection with $\mathcal{S}(v_1)$. Knowing $\mathcal{S}^*(v_{r+g+1})$, we can identify vertex v_{r+g+2} as follows. Consider the last time the variable $S_{v_{r+g+1}}$ was modified during the execution of the algorithm before the M-th iteration, say this is during the p_1-th iteration. Thus, during that iteration, $S_{v_{r+g+1}}$ was assigned the set $\mathcal{S}^*(v_{r+g+1})$, and this did not trigger the existence of a bad path. Then, the children $w_1, \ldots, w_k$ of v_{r+g+1} were inspected one by one in order, until a problematic child w_j was found. This problematic child w_j is vertex v_{r+g+2}. This process is completely deterministic, thus we can simulate it. (Indeed, we know the whole sublist assignment for vertices of A at the beginning of the p_1-th iteration, and we know the sublist that

was sampled for v_{r+g+1} during that iteration, namely, $\mathcal{S}^*(v_{r+g+1})$). Thus, for each vertex $v \in \mathrm{UP}(w_1) \cup \cdots \cup \mathrm{UP}(w_{j-1})$, we can figure out how the sublist S_v was set during that iteration, and this is exactly the value of $\mathcal{S}^*(v)$, since those sublists have not been modified afterwards prior to iteration M. Notice also that $\mathcal{S}^*(v) = \texttt{undefined}$ for all $v \in \mathrm{UP}(w_{j+1}) \cup \cdots \cup \mathrm{UP}(w_r)$, so it only remains to determine $\mathcal{S}^*(v)$ for $v \in \mathrm{UP}(v_{r+g+2})$.

We can iterate this argument and discover step by step vertices $v_{r+g+2}, \ldots, v_{2r+g}$ and the missing entries of $\mathcal{S}^*$. We spell out the general argument now, for the sake of completeness (the reader who is already convinced that it can be done is invited to skip this paragraph). For each index $i = 2, \ldots, r-1$, we proceed as follows. The set $\mathcal{S}^*(v_{r+g+i})$ is the γ_i-th ℓ-sublist of $L_{v_{r+g+i}}$ having a non-empty intersection with $\mathcal{S}(v_i)$, and is thus determined by $(D, \mathcal{S}, B, \Gamma)$. Knowing $\mathcal{S}^*(v_{r+g+i})$, we now identify vertex $v_{r+g+i+1}$. Say $S_{v_{r+g+i}}$ was modified for the last time during the p_i-th iteration of the algorithm before the M-th iteration. During that iteration, variable $S_{v_{r+g+i}}$ was assigned the set $\mathcal{S}^*(v_{r+g+i})$, and this did not result in the existence of any bad path. Next, the children $w_1, \ldots, w_k$ of v_{r+g+i} were inspected until a problematic child w_j was found, which is vertex $v_{r+g+i+1}$. We know the whole sublist assignment for vertices of A at the beginning of the p_i-th iteration; indeed, this is exactly the one at the *end* of the p_{i-1}-th iteration (as is easily checked), which we already know. We also know the sublist that was sampled for vertex v_{r+g+i} during the p_i-th iteration. Hence, we can simulate the execution of lines 12–16 of the algorithm for the p_i-th iteration. This implies that, for each vertex $v \in \mathrm{UP}(w_1) \cup \cdots \cup \mathrm{UP}(w_{j-1})$, we can determine how the sublist S_v was set during that iteration, and that sublist is precisely the set $\mathcal{S}^*(v)$. Also, we have that $\mathcal{S}^*(v) = \texttt{undefined}$ for all $v \in \mathrm{UP}(w_{j+1}) \cup \cdots \cup \mathrm{UP}(w_r)$.

This way we completely determined $\mathcal{S}^*(v)$ for all vertices $v \in \mathrm{UP}(v_{r+g+i}) - \mathrm{UP}(v_{r+g+i+1})$ for each $i \in [r-1]$. For all other vertices v of A, we have $\mathcal{S}^*(v) = \mathcal{S}(v)$. Thus, $\mathcal{S}^*$ is completely determined by $(D, \mathcal{S}, B, \Gamma)$, and hence so is the $(M-1)$-th log $(D^*, \mathcal{S}^*, B^*, \Gamma^*)$, as claimed.

Equipped with the knowledge of $\mathcal{S}^*$, we may now finish the proof in a manner similar to the $r = 1$ case. Let X denote the r_M-th ℓ-sublist X of $L_{v_{2r+g}}$. The sublist X is the γ_r-th ℓ-sublist of $L_{v_{2r+g}}$ having a non-empty intersection with $\mathcal{S}(v_r)$, which is thus determined by the log $(D, \mathcal{S}, B, \Gamma)$. Hence, we can obtain r_M from $(D, \mathcal{S}, B, \Gamma)$. Moreover, as we have seen, $(D^*, \mathcal{S}^*, B^*, \Gamma^*)$ is completely determined by the log $(D, \mathcal{S}, B, \Gamma)$. By induction, the sequence $r_1, \ldots, r_{M-1}$ is uniquely determined by $(D^*, \mathcal{S}^*, B^*, \Gamma^*)$, and therefore it is uniquely determined by $(D, \mathcal{S}, B, \Gamma)$. This concludes the proof. □

We remark that no effort has been made to optimize the bound of $32\ell^3+1$ in Lemma 6.

Acknowledgements

We thank the two anonymous referees for their insightful comments, which greatly helped us improve the exposition.

References

[1] N. Alon, J. Grytczuk, M. Hałuszczak, and O. Riordan. Nonrepetitive colorings of graphs. *Random Structures Algorithms*, 21(3-4):336–346, 2002.

[2] J. Berstel. Axel Thue's papers on repetitions in words: a translation. *Publications du LaCIM*, 20, 1995. Université du Québec à Montréal.

[3] B. Brešar, J. Grytczuk, S. Klavžar, S. Niwczyk, and I. Peterin. Nonrepetitive colorings of trees. *Discrete Math.*, 307(2):163–172, 2007.

[4] R. Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, Heidelberg, fourth edition, 2010.

[5] V. Dujmović, F. Frati, G. Joret, and D. R. Wood. Nonrepetitive colourings of planar graphs with $O(\log n)$ colours. *Electronic J. Combinatorics*, 20(1):P51, 2013. arXiv:1202.1569.

[6] V. Dujmović, G. Joret, J. Kozik, and D. R. Wood. Nonrepetitive colouring via entropy compression. to appear in *Combinatorica*, arXiv:1112.5524.

[7] F. Fiorenzi, P. Ochem, P. Ossona de Mendez, and X. Zhu. Thue choosability of trees. *Discrete Appl. Math.*, 159(17):2045–2049, 2011.

[8] J. Grytczuk. Nonrepetitive graph coloring. In *Graph Theory in Paris*, Trends in Mathematics, pages 209–218. Birkhauser, 2007.

[9] J. Grytczuk. Pattern avoidance on graphs. *Discrete Math.*, 307(11-12):1341–1346, 2007.

[10] J. Grytczuk, J. Kozik, and P. Micek. A new approach to nonrepetitive sequences. *Random Structures Algorithms*, 42(2):214–225, 2013. arXiv:1103.3809.

[11] J. Grytczuk, J. Przybyło, and X. Zhu. Nonrepetitive list colourings of paths. *Random Structures Algorithms*, 38(1-2):162–173, 2011.

[12] J. Harant and S. Jendrol'. Nonrepetitive vertex colorings of graphs. *Discrete Math.*, 312(2):374–380, 2012.

[13] J. Kozik and P. Micek. Nonrepetitive choice number of trees. *SIAM J. Discrete Math.*, 27(1):436–446, 2013. arXiv:1207.5155.

[14] A. Kündgen and M. J. Pelsmajer. Nonrepetitive colorings of graphs of bounded tree-width. *Discrete Math.*, 308(19):4473–4478, 2008.

[15] R. A. Moser and G. Tardos. A constructive proof of the general Lovász Local Lemma. *J. ACM*, 57(2):Art. 11, 15, 2010. arXiv:0903.0544.

[16] A. Thue. Über unendliche Zeichenreichen. *Norske Vid. Selsk. Skr., I Mat. Nat. Kl., Christiania*, 1–22, 1906.